

Optimal Observation Time Points in Stochastic Chemical Kinetics

Charalampos Kyriakopoulos¹ and Verena Wolf¹

Saarland University, Saarbrücken, Germany

Abstract. Wet-lab experiments, in which the dynamics within living cells are observed, are usually costly and time consuming. This is particularly true if single-cell measurements are obtained using experimental techniques such as flow-cytometry or fluorescence microscopy. It is therefore important to optimize experiments with respect to the information they provide about the system. In this paper we make a priori predictions of the amount of information that can be obtained from measurements. We focus on the case where the measurements are made to estimate parameters of a stochastic model of the underlying biochemical reactions. We propose a numerical scheme to approximate the Fisher information of future experiments at different observation time points and determine optimal observation time points. To illustrate the usefulness of our approach, we apply our method to two interesting case studies.

1 Introduction

The successful calibration of mathematical models of biological processes is usually achieved by the interplay between computer simulations and wet-lab experiments. While both approaches are typically very time consuming, wet-lab experiments are costly compared to computer simulations, which is particularly true for single-molecule techniques such as flow cytometry or fluorescence microscopy. To keep the effort and costs as low as possible and to employ wet-lab resources to have maximal gain, it is possible to run computer simulations before measurements have been done in order to maximize the amount of information provided by the measurements. If the plan is to use the measurements for the estimation of unknown model parameters, then it is common practice to approach the optimal experimental design problem by considering the Fisher information. The Fisher information provides an approximation of the accuracy of parameter estimates and if the information is maximal so is the (approximated) accuracy of the estimators. While classical parameter estimation techniques compute the Fisher information of parameter values estimated based on certain observations of the system, it is also possible to do this computation before any observations are made. Thus, it is possible to compare the amount of information that a certain hypothetical experiment provides w.r.t. an unknown parameter.

Here, we focus on stochastic models of biochemical reaction networks, for which the optimal experimental design problem has been addressed rarely in the past. We assume that the kinetic constants of the chemical reactions have to be estimated and that flow cytometry measurements are possible. Our goal is to find the optimal times at which measurements should be made to maximize the amount of information provided

by the observations. Note that such observations are not correlated over time, thus, the results are independent single-cell measurements. We do not optimize over other experiment design criteria such as the choice between the chemical species to observe. The reason is that the latter problem is, compared to the optimal observation time points problem, much simpler since we only have to compare the Fisher information for all possible combinations of observed species. Finding optimal observation times, however, is a challenging problem since the computation of the Fisher information relies on a transient solution of the model. It is therefore very costly and can usually not be done for a very large number of time points. Thus, a sophisticated numerical procedure is necessary to efficiently determine optimal observation times. Another problem is that the Fisher information depends not only on the time points the observations are made, but also on the unknown parameters value. We therefore assume that prior knowledge about the unknown parameters is available in the form of a prior distribution. Given such a prior and a stochastic model with unknown kinetic constants, we determine those observation time points at which the expected Fisher information is maximal.

There are two previous approaches to the optimal experimental design problem for stochastic chemical kinetics. Both of them provide approximations of the Fisher information and focus on systems for which a direct numerical computation of the transient solution is too expensive. Komorowski et al. propose an approach that is based on the linear noise approximation, which assumes that molecules are present in sufficiently high copy numbers [7]. However, many systems involve species present in small copy numbers leading to highly skewed distributions [15]. In such cases the linear noise approximation yields poor approximations and the underlying master equation has to be solved directly. Ruess et al. propose an approximation of the Fisher information based on the moments of the underlying probability distribution. Assuming that the number of observed cells is large, they derive an expression for the Fisher information that only involves moments up to order four [15]. This derivation is based on the fact that sample mean and sample variance of the observations are approximately normally distributed. However, it turns out that the information provided by the sample mean and the sample variance do not suffice for the characterization of skewed or bimodal distributions. In this paper we use a method based on the direct approximation of the underlying probability distribution since we found approximation errors up to 40% when using the approach of Ruess et al. We combine this approximation with a gradient-descent based optimization scheme in which we, after sampling from the prior distribution, have to solve the underlying master equation only once over time. The latter is achieved by a direct numerical approximation where the large state space of the model is truncated dynamically. This approach is a modification of an earlier method developed for the estimation of parameters [1].

After introducing the stochastic model in Section 2, we give the mathematical background for the estimation of unknown parameters and explain how the Fisher information and its time-derivatives are computed in Section 3. In Section 4 we present the optimal experiment design problem and propose a numerical method for finding optimal observation times. Finally, we report on experimental results for two reaction networks (Section 5) and conclude with a discussion of the results in Section 6.

2 Discrete-state Stochastic Model

According to Gillespie’s theory of stochastic chemical kinetics, a well-stirred mixture of n molecular species in a volume with fixed size and fixed temperature can be represented as a continuous-time Markov chain $\{\mathbf{X}_t, t \geq 0\}$ [5]. The random vector \mathbf{X}_t describes the chemical populations at time t , i.e., the i -th entry is the number of molecules of type $i \in \{1, \dots, n\}$ at time t . Thus, the state space of \mathbf{X} is $\mathbb{Z}_+^n = \{0, 1, \dots\}^n$. The state changes of \mathbf{X} are triggered by the occurrences of chemical reactions, which are of m different types. For $j \in \{1, \dots, m\}$ let the row vector $\mathbf{v}_j \in \mathbb{Z}^n$ be the nonzero *change vector* of the j -th reaction type, that is, $\mathbf{v}_j = \mathbf{v}_j^- + \mathbf{v}_j^+$ where \mathbf{v}_j^- contains only non-positive entries, which specify how many molecules of each species are consumed (*reactants*) if an instance of the reaction occurs. The vector \mathbf{v}_j^+ contains only non-negative entries, which specify how many molecules of each species are produced (*products*). Thus, if $\mathbf{X}_t = \mathbf{x}$ for some $\mathbf{x} \in \mathbb{Z}_+^n$ with $\mathbf{x} + \mathbf{v}_j^-$ being non-negative, then $\mathbf{X}_{t+dt} = \mathbf{x} + \mathbf{v}_j$ is the state of the system after the occurrence of the j -th reaction within the infinitesimal time interval $[t, t + dt)$.

Each reaction type has an associated *propensity function*, denoted by $\alpha_1, \dots, \alpha_m$, which is such that $\alpha_j(\mathbf{x}) \cdot dt$ is the probability that, given $\mathbf{X}_t = \mathbf{x}$, one instance of the j -th reaction occurs within $[t, t + dt)$. Often the value $\alpha_j(\mathbf{x})$ is chosen proportional to the number of distinct reactant combinations in state \mathbf{x} , known as the law of mass action. However, for many reactions the proportionality constant θ_j is unknown and has to be estimated based on measurements. For instance, if we have two (distinct) reactants (i.e. $\mathbf{v}_j^- = -\mathbf{e}_i - \mathbf{e}_\ell$) then $\alpha_j(\mathbf{x}) = \theta_j \cdot x_i \cdot x_\ell$ where x_i and x_ℓ are the corresponding entries of \mathbf{x} , $i \neq \ell$, $\theta_j > 0$, and \mathbf{e}_i is the vector with the i -th entry 1 and all other entries 0. In the sequel we do not restrict the form of α_j but only assume that its derivative w.r.t. some unknown parameter θ_j exists. Sometimes we will make the dependence of α_j on θ_j explicit by writing $\alpha_j(\mathbf{x}, \theta_j)$ instead of $\alpha_j(\mathbf{x})$.

Example 1. We consider a simple crystallization process that involves four chemical species, namely A, B, C and D. Thus, the entries of the random vector \mathbf{X}_t give the numbers of molecules of types A, B, C and D at time t . The two possible reactions are $2A \rightarrow B$ and $A + C \rightarrow D$. Thus, $\mathbf{v}_1 = (-2, 1, 0, 0)$, $\mathbf{v}_2 = (-1, 0, -1, 1)$. For a state $\mathbf{x} = (x_A, x_B, x_C, x_D)$, the propensity functions are $\alpha_1(\mathbf{x}) = \theta_1 \cdot \binom{x_A}{2}$ and $\alpha_2(\mathbf{x}) = \theta_2 \cdot x_A \cdot x_C$. Note that given an initial state \mathbf{x}_0 the set of reachable states is a finite subset of \mathbb{N}^4 .

In general, the reaction rate constants θ_j refer to the probability that a randomly selected pair of reactants collides and undergoes the j -th chemical reaction. It depends on the volume and the temperature of the system as well as on the microphysical properties of the reactant species. Since reactions of higher order (requiring more than two reactants) are usually the result of several successive lower order reactions, we do not consider the case of more than two reactants.

The Chemical Master Equation. For $\mathbf{x} \in \mathbb{Z}_+^n$ and $t \geq 0$, let $p_t(\mathbf{x})$ denote the probability $Pr(\mathbf{X}_t = \mathbf{x})$. Given $\mathbf{v}_1^-, \dots, \mathbf{v}_m^-, \mathbf{v}_1^+, \dots, \mathbf{v}_m^+, \alpha_1, \dots, \alpha_m$, and some initial distribution \mathbf{p}_0 , the Markov chain \mathbf{X} is uniquely specified and its evolution is given

by the chemical master equation (CME)

$$\frac{d}{dt}p_t(\mathbf{x}) = \sum_{j:\mathbf{x}-\mathbf{v}_j^- \geq 0} \alpha_j(\mathbf{x}-\mathbf{v}_j)p_t(\mathbf{x}-\mathbf{v}_j) - \alpha_j(\mathbf{x})p_t(\mathbf{x}). \quad (1)$$

If we use \mathbf{p}_t to denote the row vector with entries $p_t(\mathbf{x})$, then the vector form of the CME becomes

$$\frac{d}{dt}\mathbf{p}_t = \mathbf{p}_t Q, \quad (2)$$

where Q is the infinitesimal generator matrix of \mathbf{X} with $Q(\mathbf{x}, \mathbf{y}) = \alpha_j(\mathbf{x})$ if $\mathbf{y} = \mathbf{x} + \mathbf{v}_j$ and $\mathbf{x} + \mathbf{v}_j^- \geq 0$. Note that, in order to simplify our presentation, we assume here that all vectors \mathbf{v}_j are distinct. All remaining entries of Q are zero except for the diagonal entries which are equal to the negative row sum. The ordinary first-order differential equation in (2) is a direct consequence of the Kolmogorov forward equation. Since \mathbf{X} is a regular Markov process, (2) has the general solution $\mathbf{p}_t = \mathbf{p}_0 \cdot e^{Qt}$, where e^A is the matrix exponential of a matrix A . If the state space of X is infinite, then we can only compute approximations of \mathbf{p}_t . But even if Q is finite, its size is often large because it grows exponentially with the number of state variables. Therefore standard numerical solution techniques for systems of first-order linear equations of the form of (2) are infeasible. The reason is that the number of nonzero entries in Q often exceeds the available memory capacity for systems of realistic size. If the populations of all species remain small (at most a few hundreds) then the CME can be efficiently approximated using projection methods [6,12,4] or fast uniformization methods [10,16]. The idea of these methods is to avoid an exhaustive state space exploration and, depending on a certain time interval, restrict the analysis of the system to a subset of states.

Here, we are also interested in the partial derivatives of \mathbf{p}_t w.r.t. the reaction rate constants $\theta = (\theta_1, \dots, \theta_m)$. In the sequel we will write $\mathbf{p}_t(\theta)$ instead of \mathbf{p}_t to make the dependency on θ explicit and the entry of $\mathbf{p}_t(\theta)$ that corresponds to state \mathbf{x} will be denoted by $p_t(\mathbf{x}; \theta)$. Moreover, we define the row vectors $\mathbf{s}_t^j(\theta)$ as the derivative of $\mathbf{p}_t(\theta)$ w.r.t. θ_j , i.e.,

$$\mathbf{s}_t^j(\theta) = \frac{\partial \mathbf{p}_t(\theta)}{\partial \theta_j} = \lim_{\Delta h \rightarrow 0} \frac{\mathbf{p}_t(\theta + \Delta \mathbf{h}^{(j)}) - \mathbf{p}_t(\theta)}{\Delta h},$$

where the vector $\Delta \mathbf{h}^{(j)}$ is zero everywhere except for the j -th position that is equal to Δh . We denote the entry in $\mathbf{s}_t^j(\theta)$ that corresponds to state \mathbf{x} by $s_t^j(\mathbf{x}, \theta)$. Derivating (2), we find that $\mathbf{s}_t^j(\theta)$ is the unique solution of the following linear system of ODEs

$$\frac{d}{dt}\mathbf{s}_t^j(\theta) = \mathbf{s}_t^j(\theta)Q + \mathbf{p}_t(\theta)\frac{\partial}{\partial \theta_j}Q, \quad (3)$$

where $j \in \{1, \dots, m\}$. The initial condition is $s_0^j(\mathbf{x}, \theta) = 0$ for all \mathbf{x} and θ since $p_0(\mathbf{x}; \theta)$ is independent of θ_j .

3 Observations and Fisher Information

Following the notation in [14], we assume that observations of a biochemical network are made at time instances $t_1, \dots, t_R \in \mathbb{R}_{\geq 0}$, where $t_1 \leq \dots \leq t_R$. The entries

of the random vector $\mathbf{O}_{t_k} \in \mathbb{R}^n$ describe the molecule numbers observed at time t_k for $k \in \{1, \dots, R\}$. Since these observations are typically subject to measurement errors, we may assume that $\mathbf{O}_{t_k} = \mathbf{X}_{t_k} + \epsilon_{t_k}$, where the entries of the error terms ϵ_{t_k} are independent and identically normally distributed with mean zero and standard deviation σ . Note that \mathbf{X}_{t_k} is the true population vector at time t_k . Clearly, this implies that, conditional on \mathbf{X}_{t_k} , the random vector \mathbf{O}_{t_k} is independent of all other observations as well as independent of the history of \mathbf{X} before time t_k .

Let f denote the joint density of $\mathbf{O}_{t_1}, \dots, \mathbf{O}_{t_R}$ and assume that $\mathbf{o}_1, \dots, \mathbf{o}_R \in \mathbb{R}^n$. Then the likelihood of the observation sequence $\mathbf{o}_{t_1}, \dots, \mathbf{o}_{t_R}$ is

$$\begin{aligned} \mathcal{L} &= f(\mathbf{O}_{t_1} = \mathbf{o}_1, \dots, \mathbf{O}_{t_R} = \mathbf{o}_R) \\ &= \sum_{\mathbf{x}_1} \dots \sum_{\mathbf{x}_R} f(\mathbf{O}_{t_1} = \mathbf{o}_1, \dots, \mathbf{O}_{t_R} = \mathbf{o}_R \mid \mathbf{X}_{t_1} = \mathbf{x}_1, \dots, \mathbf{X}_{t_R} = \mathbf{x}_R) \quad (4) \\ &\quad Pr(\mathbf{X}_{t_1} = \mathbf{x}_1, \dots, \mathbf{X}_{t_R} = \mathbf{x}_R). \end{aligned}$$

We assume that for the unobserved process \mathbf{X} we do not know the values of the rate constants $\theta = (\theta_1, \dots, \theta_m)$ and our aim is to estimate these constants.

In the sequel, in order to keep the notation simpler, we assume no measurement errors in our observations. Nevertheless, it is straightforward to extend our proposed optimal design procedure to the case where the measurements are not exact. As shown in [1] this only introduces additional weights during the calculation of the likelihood. In this case, also, one can consider the standard deviation of the error terms, σ , as an unknown parameter and apply the proposed design. Here, though, we assume that $\mathbf{O}_{t_k} = \mathbf{X}_{t_k}$ for all k and for a concrete observation sequence $\mathbf{x}_1, \dots, \mathbf{x}_R \in \mathbb{N}^n$ the likelihood becomes

$$\mathcal{L} = Pr(\mathbf{X}_{t_1} = \mathbf{x}_1, \dots, \mathbf{X}_{t_R} = \mathbf{x}_R). \quad (5)$$

Of course, \mathcal{L} depends on the chosen rate parameters θ since the probability measure $Pr(\cdot)$ does. When necessary, we will make this dependence explicit by writing $\mathcal{L}(\theta)$ instead of \mathcal{L} . Since our observations are equal to the true state \mathbf{X}_{t_k} at time t_k , we write $\mathcal{L}(\mathbf{X}; \theta)$ where with some abuse of notation \mathbf{X} now refers to the sequence $\mathbf{X}_{t_1}, \dots, \mathbf{X}_{t_R}$. We now seek constants θ^* such that

$$\theta^* = \operatorname{argmax}_{\theta} \mathcal{L}(\mathbf{X}; \theta), \quad (6)$$

where the maximum is taken over all vectors θ with all components strictly positive. This optimization problem is known as the maximum likelihood problem [8]. Note that θ^* is a random variable, in the sense it depends on the (random) observations $\mathbf{X} = (\mathbf{X}_{t_1}, \dots, \mathbf{X}_{t_R})$.

The maximum likelihood estimator θ^* is known to be asymptotically normally distributed and its covariance matrix approaches the Cramér-Rao bound $\mathcal{I}_{\mathbf{X}}(\theta)^{-1}$, where $\mathcal{I}_{\mathbf{X}}(\theta)$ is the Fisher information matrix (FIM). Note that this bound is a lower bound on the estimator's covariance matrix. It is commonly used to derive confidence intervals for the estimated values of the parameters. Previous experimental results show that the variances approximated based on the FIM are close to the variances approximated based on many repetitions of the experiments and the estimation procedure [1]. Thus,

in order to have an accurate estimation of the parameters the (co-)variances should be as small as possible, i.e., $\mathcal{I}_{\mathbf{X}}(\theta)$ should be large to achieve tight confidence intervals for θ^* .

Given an observation sequence, \mathbf{X} of the process, the entry of the FIM that corresponds to the unknown parameters θ_i and θ_j , $1 \leq i, j \leq m$ is defined as

$$(\mathcal{I}_{\mathbf{X}}(\theta))_{i,j} = \mathbb{E}_{\mathbf{X}} \left[\left(\frac{\partial}{\partial \theta_i} \log \mathcal{L}(\mathbf{X}; \theta) \right) \left(\frac{\partial}{\partial \theta_j} \log \mathcal{L}(\mathbf{X}; \theta) \right) \right]. \quad (7)$$

Note that the expectation is taken w.r.t. the observation sequence $\mathbf{X} = (\mathbf{X}_{t_1}, \dots, \mathbf{X}_{t_R})$. Under certain (mild) regularity conditions that hold for the likelihoods we consider here, Eq. (7) can be equivalently written as [17]

$$\begin{aligned} (\mathcal{I}_{\mathbf{X}}(\theta))_{i,j} &= -\mathbb{E}_{\mathbf{X}} \left[\frac{\partial^2}{\partial \theta_i \partial \theta_j} \log \mathcal{L}(\mathbf{X}; \theta) \right] \\ &= -\sum_{\mathbf{x}_1, \dots, \mathbf{x}_R} Pr(\mathbf{X}_{t_1} = \mathbf{x}_1, \dots, \mathbf{X}_{t_R} = \mathbf{x}_R; \theta) \\ &\quad \frac{\partial^2}{\partial \theta_i \partial \theta_j} \log Pr(\mathbf{X}_{t_1} = \mathbf{x}_1, \dots, \mathbf{X}_{t_R} = \mathbf{x}_R; \theta). \end{aligned} \quad (8)$$

Note that if the observations $\mathbf{X}_{t_1}, \dots, \mathbf{X}_{t_R}$ are independent observations then the (i, j) th entry of the FIM is such that

$$(\mathcal{I}_{\mathbf{X}}(\theta))_{i,j} = \sum_{k=1}^R (\mathcal{I}_{\mathbf{X}_{t_k}}(\theta))_{i,j} \quad (9)$$

where $\mathcal{I}_{\mathbf{X}_{t_k}}(\theta)$ is the Fisher information matrix of a single observation \mathbf{X}_{t_k} at time t_k . The above means that the information of a sequence of R independent observations is simply the sum of the information of each. This makes the computation of the total information easier than in the dependent case since it is enough to solve the CME along with the partial derivatives $s_t^j(\theta)$, for all j , until time point t_R . This can be easily seen by exploiting (7) for \mathbf{X}_{t_k} .

$$\begin{aligned} (\mathcal{I}_{\mathbf{X}_{t_k}}(\theta))_{i,j} &= \mathbb{E}_{\mathbf{X}_{t_k}} \left[\left(\frac{\partial}{\partial \theta_i} \log Pr(\mathbf{X}_{t_k}; \theta) \right) \left(\frac{\partial}{\partial \theta_j} \log Pr(\mathbf{X}_{t_k}; \theta) \right) \right] \\ &= \sum_{\mathbf{x}_k} \left(\frac{\partial}{\partial \theta_i} \log p_{t_k}(\mathbf{x}_k; \theta) \right) \left(\frac{\partial}{\partial \theta_j} \log p_{t_k}(\mathbf{x}_k; \theta) \right) p_{t_k}(\mathbf{x}_k; \theta) \\ &= \sum_{\mathbf{x}_k} \frac{1}{p_{t_k}(\mathbf{x}_k; \theta)} \frac{\partial}{\partial \theta_i} p_{t_k}(\mathbf{x}_k; \theta) \frac{\partial}{\partial \theta_j} p_{t_k}(\mathbf{x}_k; \theta), \end{aligned} \quad (10)$$

where the sums run over all possible states \mathbf{x}_k that can be observed at time t_k . Using the notation of the previous section, we get

$$(\mathcal{I}_{\mathbf{X}_{t_k}}(\theta))_{i,j} = \sum_{\mathbf{x}_k} \frac{s_{t_k}^i(\mathbf{x}_k; \theta) s_{t_k}^j(\mathbf{x}_k; \theta)}{p_{t_k}(\mathbf{x}_k; \theta)}. \quad (11)$$

Derivating (11) we can also compute the derivative of the FIM w.r.t. the time point t_k .

$$\begin{aligned} \frac{\partial}{\partial t} (\mathcal{I}_{\mathbf{x}_{t_k}}(\theta))_{i,j} &= \sum_{\mathbf{x}_k} \frac{\frac{\partial}{\partial t} s_{t_k}^i(\mathbf{x}_k; \theta) s_{t_k}^j(\mathbf{x}_k; \theta) + s_{t_k}^i(\mathbf{x}_k; \theta) \frac{\partial}{\partial t} s_{t_k}^j(\mathbf{x}_k; \theta)}{p_{t_k}(\mathbf{x}_k; \theta)} \\ &\quad - \sum_{\mathbf{x}_k} \frac{s_{t_k}^i(\mathbf{x}_k; \theta) s_{t_k}^j(\mathbf{x}_k; \theta) \frac{\partial}{\partial t} p_{t_k}(\mathbf{x}_k; \theta)}{p_{t_k}(\mathbf{x}_k; \theta)^2}. \end{aligned} \quad (12)$$

The time derivative of the FIM is particularly useful for an efficient gradient-based optimization scheme to find the time points that provide the maximum information which is the main goal of the next section.

4 Optimal Observation Time Points

Assume now that an experiment is planned where the system is observed at time points t_1, \dots, t_R and that the observations will be independent (e.g. if the chosen measurement technique is flow cytometry). We want to find the optimal time points to take our observations, i.e. those which yield the maximum information for the system parameters θ . However, the intrinsic problem in experiment design is that the parameter values θ are, of course, unknown before the experiment is set up. Here, the chosen approach to overcome this obstacle is to search for the observation time points that maximize the determinant of the *expected* FIM when one assumes a prior distribution for the unknown parameters as suggested, for instance, in [11]. In other words our goal is to find

$$\mathbf{t}^* = \underset{\mathbf{t}=(t_1, \dots, t_R)}{\operatorname{argmax}} \det(\mathbb{E}_\theta[\mathcal{I}(\mathbf{t}, \theta)]). \quad (13)$$

This criterion is known to be robust because it incorporates a prior belief for θ . Note that in the sequel we write $\mathcal{I}(\mathbf{t}, \theta)$ instead of $\mathcal{I}_{\mathbf{X}}(\theta)$ to explicitly indicate that the information is a function of the sequence of time points, \mathbf{t} , at which the observations are to be taken. Also, remember that in the previous section we showed that $\mathcal{I}(\mathbf{t}, \theta) = \sum_{k=1}^R \mathcal{I}(t_k, \theta)$, when we assume that the observations are independent.

It is worth mentioning that sometimes not all reaction rate constants are of interest, in which case we partition the parameter vector $\theta = [\theta_{(1)}, \theta_{(2)}]$ in parameters of interest $\theta_{(1)}$ and nuisance parameters $\theta_{(2)}$. Then, for a specific time point t , $\mathcal{I}(t, \theta)$ is replaced by the matrix $\mathcal{I}_s(t, \theta)$ which is defined as

$$\begin{aligned} \mathcal{I}_s(t, \theta) &= \mathcal{I}_{11}(t, \theta) - \mathcal{I}_{12}(t, \theta) \mathcal{I}_{22}^{-1}(t, \theta) \mathcal{I}_{12}^\top(t, \theta), \text{ where} \\ \mathcal{I}(t, \theta) &= \begin{bmatrix} \mathcal{I}_{11}(t, \theta) & \mathcal{I}_{12}(t, \theta) \\ \mathcal{I}_{12}^\top(t, \theta) & \mathcal{I}_{22}(t, \theta) \end{bmatrix}. \end{aligned} \quad (14)$$

Here, subscript s indicates that we only consider derivatives of the parameters of interest, i.e., the matrices $\mathcal{I}_{11}(t, \theta)$, $\mathcal{I}_{22}(t, \theta)$ contain information about the variance of $\theta_{(1)}$ and $\theta_{(2)}$, respectively, while $\mathcal{I}_{12}(t, \theta)$ approximates the covariance matrix of $\theta_{(1)}$ and $\theta_{(2)}$. Hence, given a prior of θ and having computed the matrices $\mathcal{I}(t_k, \theta)$ for all θ and all k , it is straightforward to compute the matrix $\mathbb{E}_\theta[\mathcal{I}_s(\mathbf{t}, \theta)]$ and then $\det(\mathbb{E}_\theta[\mathcal{I}_s(\mathbf{t}, \theta)])$.

Similarly, if we have, in addition, $\frac{\partial}{\partial t} \mathcal{I}(t, \theta)$ for all θ and all $t \in \{t_1, \dots, t_R\}$ we can also compute $\frac{\partial}{\partial t} \det(\mathbb{E}_\theta[\mathcal{I}_s(\mathbf{t}, \theta)])$ as follows. From Jacobi's formula it holds that for any square matrix A

$$\frac{\partial}{\partial t} \det(A) = \text{tr}(\text{adj}(A)) \frac{\partial}{\partial t} A, \quad (15)$$

where $\text{tr}(A)$ is defined as the sum of the elements of the main diagonal of A and $\text{adj}(A)$ is the adjoint matrix of A . Here, $A = \mathbb{E}_\theta[\mathcal{I}_s(\mathbf{t}, \theta)]$ and, derivating Eq. (14), we can compute $\frac{\partial}{\partial t} \mathbb{E}_\theta[\mathcal{I}_s(\mathbf{t}, \theta)]$ by exploiting known matrix calculus identities.

$$\begin{aligned} \frac{\partial}{\partial t} \mathbb{E}_\theta[\mathcal{I}_s(\mathbf{t}, \theta)] &= \mathbb{E}_\theta \left[\frac{\partial}{\partial t} \mathcal{I}_{11}(\mathbf{t}, \theta) - \frac{\partial}{\partial t} \mathcal{I}_{12}(\mathbf{t}, \theta) \mathcal{I}_{22}^{-1}(\mathbf{t}, \theta) \mathcal{I}_{12}^\top(\mathbf{t}, \theta) \right. \\ &\quad \left. - \mathcal{I}_{12}(\mathbf{t}, \theta) \frac{\partial}{\partial t} \mathcal{I}_{22}^{-1}(\mathbf{t}, \theta) \mathcal{I}_{12}^\top(\mathbf{t}, \theta) \right. \\ &\quad \left. - \mathcal{I}_{12}(\mathbf{t}, \theta) \mathcal{I}_{22}^{-1}(\mathbf{t}, \theta) \frac{\partial}{\partial t} \mathcal{I}_{12}^\top(\mathbf{t}, \theta) \right], \end{aligned} \quad (16)$$

where the derivative of the inverse is computed as

$$\frac{\partial}{\partial t} \mathcal{I}_{22}^{-1}(\mathbf{t}, \theta) = -\mathcal{I}_{22}^{-1}(\mathbf{t}, \theta) \frac{\partial}{\partial t} \mathcal{I}_{22}(\mathbf{t}, \theta) \mathcal{I}_{22}^{-1}(\mathbf{t}, \theta).$$

The main computational effort in the search for the optimal experiment is that for finding \mathbf{t}^* as defined in Eq. (13), we have to solve the CME after sampling from the prior of θ . It is generally not possible to find \mathbf{t}^* from a single solution of the CME since we need to average over all possible values for $\theta = (\theta_1, \dots, \theta_m)$. Therefore, we propose the following gradient descent based procedure to approximate local maxima of the determinant of the expected FIM:

1. Given a prior μ_θ for the distribution of θ , sample values $\theta^{(1)}, \dots, \theta^{(N)} \sim \mu_\theta$.
2. Choose a sequence $\mathbf{t}_{\text{next}} = (t_1, \dots, t_R)$ and for each sample of $\theta^{(i)}$ compute $\mathcal{I}_s(\mathbf{t}_{\text{next}}, \theta^{(i)})$ and $\frac{\partial}{\partial t} \mathcal{I}_s(\mathbf{t}_{\text{next}}, \theta^{(i)})$ for $t \in \{t_1, \dots, t_R\}$.
3. Return approximations of $\det(\mathbb{E}_\theta[\mathcal{I}_s(\mathbf{t}, \theta)])$ and $\frac{\partial}{\partial t} \det(\mathbb{E}_\theta[\mathcal{I}_s(\mathbf{t}, \theta)])$ by averaging over the results for $i = 1, \dots, N$.
4. Following the gradient choose $\mathbf{t}_{\text{next}} = (t'_1, \dots, t'_R)$ and repeat from 2 until you find a local maximum.

Clearly, an approximation of the global maximum is found by starting the local gradient based search from multiple initial points. The initial points used in the second step can be chosen randomly or according to some heuristics as it is usual for global optimization methods. A technical but computationally important detail is that there is no need to solve the CME for every sequence of time points that is considered in the optimization algorithm. We can solve the CME once and keep the values of the Fisher information matrix and its derivatives over time and recall it for every new sequence of time points (up to a chosen discretization). Consequently, the total computational effort of the above optimization procedure is to solve once the CME until time point equal to the maximal value of t_R encountered during the optimization.

Certainly, the above optimality criterion is by no means the only possible choice. A slightly different approach, for instance, would be to maximize the expected determinant of FIM [13]. Alternatively, if no prior is available, it is also possible to consider

$\max_t \min_\theta \det(\mathcal{I}_s(t, \theta))$ to make sure that for any choice of θ the experiment provides maximal information. Here, we assume that a prior is available because, most probably, one experiment has been already done in order to acquire some prior knowledge for the parameters. At last, an additional advantage of a robust optimal design is that in case there is the option to perform more than two experiments in total, the above procedure can be used in iterations alternating between experiments and the update of the prior of θ via parameter estimation from the real observations.

5 Experimental Results

We consider two biochemical reaction networks to which we apply our experiment design procedure, namely the crystallization model, described in Example 1, and the so-called exclusive switch model [9]. The crystallization model is a very simple example because it has a finite state space and the CME can be integrated directly if the initial molecule numbers are not particularly large. However, for large initial conditions a transient solution is only possible if the state space is dynamically truncated as suggested in [10] (see also [2]). The second example is infinite in two dimensions and its distribution is bimodal. Thus, for this system we solved the CME using a dynamic truncation of the state space whenever necessary. We chose the truncation threshold e^{-20} and the dynamic truncation of the states are based on the ratio in Eq. (11), rather than only on the value of $p_t(\mathbf{x}; \theta)$. More precisely, a state \mathbf{x} is considered as significant at time t whenever this ratio is greater than the truncation threshold while in [10] it is only the current probability that determines whether a state is considered or not at time t .

5.1 Crystallization





Fig. 1. Approximated Fisher information of the crystallization example for different values of θ_2 .

Recall the simple crystallization process given by the two chemical reactions



Let us initially assume that the value of $\theta_1 = 4$ is known and we have a prior for $\theta_2 \sim \mathbb{U}(0.05, 0.1)$ where $\mathbb{U}(a, b)$ refers to the continuous uniform distribution between a and b . For the initial state $(x_A, x_B, x_C, x_D) = (4, 0, 2, 0)$ the number of reachable states is 7 and no sophisticated truncation method is necessary to integrate the CME. Our goal is to find the best time points to take an observation for estimating θ_2 . In this case, the Fisher information is just a scalar and in

Figure 1 it is shown as a function of both θ_2 and time. The different colors correspond to different values of $\mathcal{I}(t, \theta_2)$ (see colorbar). The time point for the maximum expected

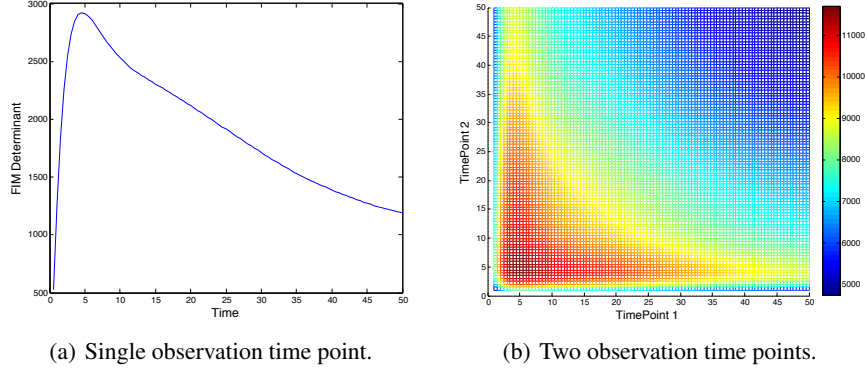


Fig. 2. Approximated determinant of the expected FIM over time for the crystallization example.

information is $t^* = 20.64$. Note that in case of one unknown parameter the extension to the R time points optimization problem is straightforward because of the linearity of (9). The solution simply consists of R replications of t^* .

Next, we assume that both, θ_1 and θ_2 are unknown parameters of interest. Then the Fisher information is a 2×2 matrix. In Figure 2(a) the evolution of the determinant of the expected FIM is shown for $\theta_1 \sim \mathbb{U}(0.05, 0.5)$ and $\theta_2 \sim \mathbb{U}(0.01, 0.1)$. From the plot it is clear that if there is only one observation possible this has to be done quite early in time. Indeed, our optimal design scheme returns that the optimal time point to take an observation is at $t = 4.625$ assuming that both parameters have to be estimated.

Now, we consider the case that we are able to take a second measurement $t_2 \geq t_1$. The determinant of the average FIM is shown in Figure 2(b). The plot suggests that one obtains the maximum information for the estimation of both parameters if both observations take place early in the experiment. Our optimization procedure returned the optimal time points $t_1 = t_2 = 4.625$. One notes that combining an early observation with a later seems also a good choice. On the other hand, taking two late observations provides significantly less information leading possibly to parameter identifiability problems.

Generalizing the above observations for the case of R time points we verify from our experiments until $R = 5$ that for this particular model and this choice of parameters the optimal experiment persistently consists of taking as many as possible observations at the same early time point. We expect the same to hold for any $R \geq 2$.

5.2 Exclusive Switch

The exclusive switch is a gene regulatory network that consists of two genes with a common promoter region as shown in Figure 3. The system involves five chemical species DNA , P_1 , P_2 , $DNA.P_1$, $DNA.P_2$. At each time point the system can be in one of the following three configurations: a) The promoter region is free, b) P_1 binds to the promoter region or c) P_2 binds to the promoter region. Each of the two gene

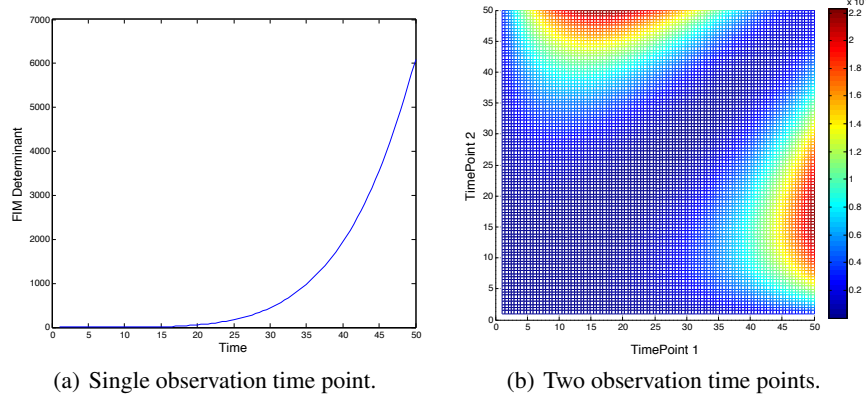


Fig. 4. Approximated determinant of the expected FIM for the exclusive switch example (case 1).

products P_1 and P_2 inhibits the expression of the other product if a molecule is bound to the promotor region. More precisely, in configuration a) (promotor region is free), molecules of both types P_1 and P_2 are produced. If a molecule of type P_1 (P_2) is bound to the promotor region (case b) and c)), only molecules of type P_1 (P_2) are produced, respectively. The chemical reactions with the corresponding constant rates are shown below for $j = \{1, 2\}$.





Fig. 3. Exclusive Switch Network

Depending on the chosen parameters, the probability distribution of the exclusive switch is bistable, i.e. most of the probability mass concentrates on two distinct regions of the state space. In particular, if binding to the promotor is likely, then these two regions correspond to the two configurations b) and c) where either the production of P_1 or the production of P_2 is inhibited.

For the purpose of our experiments we fixed the initial state of the system such that no proteins are present in the system and one DNA molecule with a free promotor region. We set up optimal experiments for the following case: The unknown parameters are the production and the degradation constants of P_1 , λ_1 and δ_1 respectively, while the rest of the parameter values are known. We assume that for $j = \{1, 2\}$

$$\lambda_1 \sim \mathbb{U}(0.01, 0.1), \delta_1 \sim \mathbb{U}(0.0001, 0.001), \\ \lambda_2 = 0.05, \delta_2 = 0.0005, \beta_j = 0.001 \text{ and } \nu_j = 0.008.$$

In Figure 4(a) the information is shown over time. From the plot it is evident that in the case of a single observation time point one should take the measurement as late as possible in the interval $[0, 50]$, if we restrict until time $t = 50$. This most probably arises from the fact that the chosen binding rates, β_1, β_2 are rather small, i.e., binding to the promotor is not so likely and there is a delay until the binding influences the dynamics of production and degradation of the corresponding proteins.

In Figure 4(b) the information of a possible experimental setup for two time points is being presented. From the plot we can observe that the most informative experiment, now, is given by two *different* time points. The second measurement is at 50 time units, as previously, but the first one should be taken at 16.75. Intuitively, this could mean that for estimating multiple parameters of this model we need to observe the process at more than one time points, if possible.

Setting up optimal experiments for R time points we observe, as in the first model, a replication of the optimal p time points, where p is the number of the unknown parameters. For $R = 3$ and $R = 4$ we get $(t_1, t_2, t_3) = (16.75, 50, 50)$ and $(t_1, t_2, t_3, t_4) = (16.75, 16.75, 50, 50)$, respectively.

6 Discussion and Future Work

Given a stochastic model of a chemical reaction network, we computed the Fisher information of different experimental designs and determined optimal observation times. The optimality criterion that we considered was the determinant of the expected Fisher information where the expectation was taken w.r.t. some prior distribution over the unknown parameters.

Our experimental results give rise to the conjecture that the n optimal time points for a system with p unknown parameters are equal replications of the p optimal time points. E.g. if we have two unknown parameters, but, say, four possible observation time points, we get only two distinct times at which the Fisher information becomes maximal w.r.t. the time points. A similar result has been proven by Box for deterministic chemical kinetics [3]. We conjecture that his result carries over to the stochastic setting. This would make experiment design for stochastic chemical kinetics a much less expensive procedure since the dimension of the optimization search space reduces to the number of unknown parameters.

Other plans for future work include that we consider equidistant observation time points and optimize the time interval between two successive observations. We also plan to consider other ways of approximating the FIM, e.g. by using moment closure techniques and exploiting the information of a sufficiently large number of moments. Additionally, we will work on the case of dependent observations and develop an iterative procedure where after each optimization step experimental results become available and the next optimal observation time is computed given these results.

References

1. A. Andreychenko, L. Mikeev, D. Spieler, and V. Wolf. Parameter identification for Markov models of biochemical reactions. In *Proc. of CAV*, volume 6806 of *Lecture Notes of Computer Science*, pages 83–98. Springer, 2011.
2. A. Andreychenko, L. Mikeev, D. Spieler, and V. Wolf. Approximate maximum likelihood estimation for stochastic chemical kinetics. *EURASIP Journal on Bioinformatics and Systems Biology*, 9, 2012.
3. G. E. P. Box and H. L. Lucas. Design of experiments in non-linear situations. *Biometrika*, 46(1/2):77–90, June 1959.
4. K. Burrage, M. Hegland, F. Macnamara, and B. Sidje. A Krylov-based finite state projection algorithm for solving the chemical master equation arising in the discrete modelling of biological systems. In *Proceedings of the Markov 150th Anniversary Conference*, pages 21–38. Boson Books, 2006.
5. D. T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *J. Phys. Chem.*, 81(25):2340–2361, 1977.
6. T. Henzinger, M. Mateescu, and V. Wolf. Sliding window abstraction for infinite Markov chains. In *Proc. CAV*, volume 5643 of *LNCS*. Springer, 2009.
7. M. Komorowski, M. J. Costa, David A. Rand, and M. P. H. Stumpf. Sensitivity, robustness, and identifiability in stochastic chemical kinetics models. *Proceedings of the National Academy of Sciences*, 108(21):8645–8650, May 2011.
8. L. Ljung. *System Identification: Theory for the User (2nd Edition)*. Prentice Hall PTR, 1998.
9. A. Loinger, A. Lipshtat, N. Q. Balaban, and O. Biham. Stochastic simulations of genetic switch systems. *Physical Review E*, 75:021904, 2007.
10. M. Mateescu, V. Wolf, F. Didier, and T.A. Henzinger. Fast adaptive uniformisation of the chemical master equation. *IET Systems Biology*, 4(6):441–452, 2010.
11. Y. Merlé and F. Mentré. Bayesian design criteria: Computation, comparison, and application to a pharmacokinetic and a pharmacodynamic model. *Journal of Pharmacokinetics and Biopharmaceutics*, 23(1):101–125, February 1995.
12. B. Munsky and M. Khammash. The finite state projection algorithm for the solution of the chemical master equation. *J. Chem. Phys.*, 124:044144, 2006.
13. L. Pronzato and E. Walter. Robust experiment design via stochastic approximation. *Mathematical Biosciences*, 75(1):103–120, 1985.
14. S. Reinker, R.M. Altman, and J. Timmer. Parameter estimation in stochastic biochemical reactions. *IEEE Proc. Syst. Biol*, 153:168–178, 2006.
15. J. Ruess, A. Miliadis-Argeitis, and J. Lygeros. Designing experiments to understand the variability in biochemical reaction networks. *Journal of The Royal Society Interface*, 10(88):20130588–20130588, August 2013.
16. R. Sidje, K. Burrage, and S. MacNamara. Inexact uniformization method for computing transient distributions of Markov chains. *SIAM J. Sci. Comput.*, 29(6):2562–2580, 2007.
17. A. van den Bos. *Parameter estimation for scientists and engineers*. Wiley-Interscience, 2007.